6.004 Computation Structures
Spring 2009

Interrupts and real time

Problem 1. A computer system has three devices whose characteristics are summarized in the following table:

Device	Service Time	Interrupt Frequency	Allowable Latency
D1	400 us	1/(800us)	400us
D2	250 us	1/(1000us)	50us
D3	100 us	1/(800us)	300us

Service time indicates how long it takes to run the interrupt handler for each device. The maximum time allowed to elapse between an interrupt request and the start of the interrupt handler is indicated by allowable latency.

A. If a program P takes 100 seconds to execute when interrupts are disabled, how long will P take to run when interrupts are enabled?

B. Can the requirements given in the table above be met using a weak priority ordering among the interrupt requests?

C. Can the requirements given in the table above be met using a strong priority ordering among the interrupt requests?

Problem 2. Surreal Time Systems is configuring a Beta for a dedicated application involving three I/O devices whose characteristics are summarized below:

Device	Interrupt Frequency	Service Time
A	1/(1000 us)	600 us
B	1/(500 us)	100 us
C	1/(1000000 us)	100000 us

Each of the three devices causes periodic interrupts as the given frequency. Each interrupt requires the service time specified for that device. When the processor is not servicing any interrupt, it runs a compute-bound user-mode (background) task L which composes limericks like

There was an old lady from Crewe
whose limericks stopped at line two.

For each of the following questions, assume all devices are requesting service at their maximum rate.

A. Assuming no interrupt priorities, what is the approximate worst-case latency seen by each device?

B. Now assume that each interrupt handler must complete execution before the next request from the same device in order to avoid losing data. To accommodate this real time constraint, the processor is enhanced with a 4-level strong priority system with priorities 0 (background), 1, 2 and 3 (highest). What priorities would you assign to each device?

C. Suppose that, in the absence of interrupts, L composes an average of 100 limericks per hour. What is its rate when all three devices are interrupting?

Problem 3. A computer system is interfaced to four devices: a printer, a disk, a keyboard, and a display. The characteristics of the devices are summarized in the following table.

Device	Interrupt service time	Interrupt frequency	response-time requirement
Printer	1000 us	1/(2000 us)	1000 us
Disk	300 us	1/(1000 us)	200 us
Keyboard	2000 us	1/(100000 us)	2000 us
Display	100 us	1/(1000 us)	200 us

A. A program P, which performs only computation (no input/output), takes 100 s to run when no input/output is being performed. How long will it take for P to run when all of the above devices are operating at their maximum speeds?

B. Suppose that the interrupt system enforces a nonpreemptive (weak) priority ordering printer > disk > keyboard > display among interrupt requests. Assuming the characteristics given in the table above, what is the maximum time that might elapse between a disk interrupt request and execution of the first instruction of its handler? Assume that the time taken for state save and context switch is negligible.

C. Can the requirements given in the table above be met using a *weak* priority ordering among the interrupt requests? If so, give such an ordering; if not, explain.

D. Can the requirements given in the table above be met using a *strong* priority ordering among the interrupt requests? If so, give such an ordering; if not, explain.

Problem 4. A computer must service three devices whose interrupting frequencies, service times, and assigned priorities are given in the table below.

Device	Service time (ms)	Maximum Frequency (1/ms)	Priority
D1	10	1/100	3 (highest)
D2	50	1/1000	2
D3	200	1/5000	1 (lowest)

A. Assuming a *strong* priority system, compute for each device the maximum time between service request and the *completion* of service for that device.

B. What percentage of the processor's time is devoted to servicing D1?

C. What percentage of the processor's time is left for noninterrupt programs?

D. Assume that if a device interrupts again before a pending interrupt on that same device has been serviced, the later interrupt is ignored (lost). Will the system outlined in the table above lose interrupts using a *strong* priority scheme (with priorities as given)?

E. Under the assumption of question (4), will the system outlined in the table above lose interrupts using a *weak* priority scheme (with priorities as given)?

Problem 5. Consider the following priority-interrupt scenario:

Task	Service time (ms)	Maximum allowed latency (ms)	Maximum Frequency (1/ms)
A	30	500	1/3000
B	20	70	1/1000

C	50	25	1/500
D	10	10	1/50

A. Can you use a weak priority scheme for the scenario outlined in the table aboved? Explain.

B. Assume that all the interrupts listed in the table above occur at their maximum frequency. What percent of the processor's time is used to handle interrupts?

C. Assume a strong priority system in which 3 is the highest priority, 0 the lowest. Assign a unique priority to each task in the table above to meet the specifications given. Show the maximum time between interrupt and *completion* of service for each of the tasks if your priority scheme is used.